Exponential node clustering at singularities for rational approximation, quadrature, and PDEs

Lloyd N. Trefethen[1] · Yuji Nakatsukasa[1] · J. A. C. Weideman[2]

Abstract

Rational approximations of functions with singularities can converge at a root-exponential rate if the poles are exponentially clustered. We begin by reviewing this effect in minimax, least-squares, and AAA approximations on intervals and complex domains, conformal mapping, and the numerical solution of Laplace, Helmholtz, and biharmonic equations by the "lightning" method. Extensive and wide-ranging numerical experiments are involved. We then present further experiments giving evidence that in all of these applications, it is advantageous to use exponential clustering whose density on a logarithmic scale is not uniform but tapers off linearly to zero near the singularity. We propose a theoretical model of the tapering effect based on the Hermite contour integral and potential theory, which suggests that tapering doubles the rate of convergence. Finally we show that related mathematics applies to the relationship between exponential (not tapered) and doubly exponential (tapered) quadrature formulas. Here it is the Gauss–Takahasi–Mori contour integral that comes into play.

1 Introduction

Analytic functions can be approximated by polynomials with exponential convergence, i.e., there exist polynomials p_n of degree n such that $\|f - p_n\| = O(\exp(-Cn))$ for some $C > 0$ as $n \to \infty$. Here $\|\cdot\|$ is the ∞-norm on an approximation domain E,

✉ Lloyd N. Trefethen
trefethen@maths.ox.ac.uk

Yuji Nakatsukasa
nakatsukasa@maths.ox.ac.uk

J. A. C. Weideman
weideman@sun.ac.za

1   Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

2   Applied Mathematics, Stellenbosch University, Private Bag X1, Matieland 7602, South Africa

which may be a closed interval of the real axis or more generally a simply connected compact set in the complex plane. This result is due to Runge [35,63] and explains the exponential convergence of many numerical methods when applied to analytic functions, including Gauss and Clenshaw–Curtis quadrature [56,58] and spectral methods for ordinary and partial differential equations [55,57]. It is also the mathematical basis of Chebfun [7].

If f is not analytic in a neighborhood of E, then Bernstein showed in 1912 that exponential convergence of polynomial approximations is impossible [3,58]. Bernstein also showed that in approximation of functions with derivative discontinuities such as $f(x) = |x|$ on $[-1, 1]$, polynomials can converge no faster than $O(n^{-1})$ [4]. Now from the beginning, going back to Chebyshev in the mid-19th century, approximation theorists had investigated approximation by rational functions as well as polynomials. Yet it was not until fifty years after these works by Bernstein that it was realized that for this problem of approximating $|x|$ on $[-1, 1]$, rational functions can achieve the much faster rate of *root-exponential convergence.* That is, there exist rational functions r_n of degree n (defined below) such that $\| f - r_n \| = O(\exp(-C\sqrt{n}))$ for some $C > 0$. This result was published by Newman [29], who also showed that faster convergence is not possible. With hindsight, it can be seen that the root-exponential effect was implicit in the results of Chebyshev's student Zolotarev nearly a century earlier [11,26,41,68], but this was not noticed.

Newman's theorem has been a great stimulus to further research in rational approximation theory [11–13,23,36,39,41,62]. It has not, however, had much impact on scientific computing until very recently with the discovery that it can be the basis of root-exponentially converging numerical methods for the solution of partial differential equations (PDEs) in domains with corner singularities [6,14–16,59]. The aim of this paper is to contribute to building the bridge between approximation theory and numerical computation.

In particular, we shall focus on the key feature that gives rational approximations their power: the exponential clustering of poles near singularities. (The zeros are also exponentially clustered, typically interlacing the poles, with the alternating pole-zero configuration serving as proxy for a branch cut.) This has been a feature of the theory since Newman's explicit construction. Our aim is, first, to show how widespread this effect is, not only with minimax approximations (i.e., optimal in the ∞-norm), the focus of most theoretical studies, but also for other kinds of approximations that may be more useful in computation. Section 2 explores this effect in a wide range of applications.

In Sect. 3 we turn to a new contribution of this paper, the observation that good approximations tend to make use of poles which, although exponentially clustered, have a density on a logarithmic scale that tapers to zero at the endpoint. Specifically, the distances of the clustered poles to the singularity appear equally spaced when the log of the distance is plotted against the square root of the index. We show experimentally that this scaling appears not just with minimax approximations but more generally.

To explain this effect, we begin with a review in Sect. 4 of the Hermite contour integral, which is the basis of the application of potential theory in approximation. We show how this leads to the idea of condenser capacity for the analysis of rational approximation of analytic functions. Section 5 then turns to functions with singularities

and proposes a model of the tapering effect. In this case the condenser is short-circuited, and it is not possible to estimate the Hermite integral by considering the ∞-norm of the factors of its integrand, but the 1-norm gives the required results. Analysis of a model problem shows how the tapered exponential clustering of poles enables better overall resolution, potentially doubling the rate of convergence. These arguments are related to those developed in the theoretical approximation theory literature by Stahl and others [39,41,43], but we believe that Sect. 5 of this paper is the first to connect this theory with numerical analysis.

Finally in Sect. 6 we turn to a different problem, the quadrature of functions with endpoint singularities on $[-1, 1]$. Here the famous methods are the exponential (tanh) and double exponential (tanh-sinh) formulas [19,21,24,25,30,50,52–54]. Making use of the link to another contour integral formula, the Gauss–Takahasi–Mori integral [10,51,56], we show that the distinction between straight and tapered exponential clustering arises here too.

Throughout the paper, R_n denotes the set of rational functions of *degree n*, that is, functions that can be written as $r(x) = p(x)/q(x)$ where p and q are polynomials of degree n. The norm $\| \cdot \|$ is the ∞-norm on E, but, as mentioned above, other measures will come into play in Sects. 5 and 6, and indeed, a theme of our discussion is that certain aspects of rational approximation are often concealed by too much focus on the ∞-norm.

The numerical experiments in this paper are a major part of the contribution; we are not aware of comparably detailed studies elsewhere in the literature. Our emphasis is on the results, not the algorithms, but our numerical methods are briefly summarized in the discussion section at the end.

2 Root-exponential convergence and exponential clustering of poles

In this section we explore the convergence of a variety of rational approximations to analytic functions with boundary branch point singularities. Our starting point is Fig. 1, which presents results for six kinds of approximations of $f(x) = \sqrt{x}$ on $[0, 1]$ by rational functions of degrees $1 \leq n \leq 20$. (By the substitution $x = t^2$, this is equivalent to Newman's problem of approximation of $|t|$ on $[-1, 1]$.) The choice of f is not special; as we shall illustrate in Figs. 2 and 4, other functions with endpoint singularities give similar results. For a related discussion of clustered poles and root exponential convergence see [20], where rational approximations associated with continued fractions are connected with optimal finite difference grids.

First, the big picture. The upper-left image of the figure shows ∞-norm errors $\|f - r_n\|$ plotted on a log scale as functions not of n but of $\sqrt{n}$. With the exception of the erratic case labeled AAA, all the curves plainly approach straight lines as $n \to \infty$: root-exponential convergence. (The shapes would be parabolas if we plotted against n.) The upper-right image shows the absolute values of the 20 poles for the approximations with $n = 20$, that is, their distances from the singularity at $x = 0$. On this logarithmic scale the poles are smoothly distributed: exponential clustering. This clustering is further shown in the lower images, for the approximation labeled

minimax, by a phase portrait [65] of the square root function (the standard branch) and its degree 20 rational approximation after an exponential change of variables.

The top four approximations have preassigned poles, making the approximation problems linear; indeed the Stenger, trapezoidal, and Newman approximations are given by explicit formulas. The AAA and minimax approximations are nonlinear, with poles determined during the computation. Although it is tempting to rank these candidates from worst at the top to best at the bottom (the minimax approximation is best by definition), this is not the point. All these approximations converge root-exponentially, and the differences among them in rates of convergence as a function of n do not exceed factors on the order of 10; the rates can in fact be further improved in most cases by introducing a scaling parameter or two. In particular, minimax and other nonlinear approximations can approximately double the rate of convergence of the linear approximations [31]. Even the slowest of these approximations can achieve accuracy 10^{-6} with degrees n not much greater than 100, whereas with polynomials one needs $n = 140,085$.

We comment now on the individual approximations of Fig. 1. The Newman approximation comes from the explicit formula presented in his four-page paper [29]. The approximation is $r(x) = \sqrt{x}(p(\sqrt{x}) - p(-\sqrt{x}))/(p(\sqrt{x}) + p(-\sqrt{x}))$, where $p(t) = \prod_{k=0}^{2n-1}(t + \xi^k)$ and $\xi = \exp(-1/\sqrt{2n})$; this can be shown to be a rational function in x of degree n. The asymptotic convergence rate is $\exp(-\sqrt{2n})$ [67]. This can be improved to approximately $\exp(-(\pi/2)\sqrt{2n})$ by defining $\xi = \exp(-(\pi/2)/\sqrt{2n})$, an example of the scaling parameters mentioned in the last paragraph (these values are conjectured to be optimal based on numerical experiments).

The trapezoidal approximation originates with Stenger's investigations of sinc functions and associated approximations [46–48]. Following p. 211 of [58], we approximate $\sqrt{x}$ by starting from the identity $\sqrt{x} = 2x/\pi \int_0^\infty (t^2 + x)^{-1} dt$, which with the change of variables $t = e^s$ becomes

$$\sqrt{x} = \frac{2x}{\pi} \int_{-\infty}^{\infty} \frac{e^s \, ds}{e^{2s} + x}. \tag{1}$$

For $n \geq 1$, we approximate this integral by an equispaced n-point trapezoidal rule with step size $h > 0$,

$$r(x) = \frac{2hx}{\pi} \sum_{k=-(n-1)/2}^{(n-1)/2} \frac{e^{kh}}{e^{2kh} + x}. \tag{2}$$

(If n is even, the values of k are half-integers.) There are n terms in the sum, so r is a rational function of degree n with simple poles at the points $p_k = -\exp(2kh)$. Two sources of error make $r(x)$ differ from $\sqrt{x}$. The termination of the sum at $n < \infty$ introduces an error of the order of $\exp(-nh/2)$, and the finite step size introduces an error on the order of $\exp(-\pi^2/h)$, since the integrand is analytic in the strip around the real s-axis of half-width $\pi/2$ [61, Theorem 5.1]. Balancing these errors gives the optimal step size $h \approx \pi\sqrt{2/n}$ and approximation error $\|r - \sqrt{x}\| \approx \exp(-\pi\sqrt{n/2})$. Note that the poles for this approximation cluster at ∞ as well as at 0, and indeed, it converges root-exponentially not just on [0, 1] but on any interval [0, L] with $L > 0$.
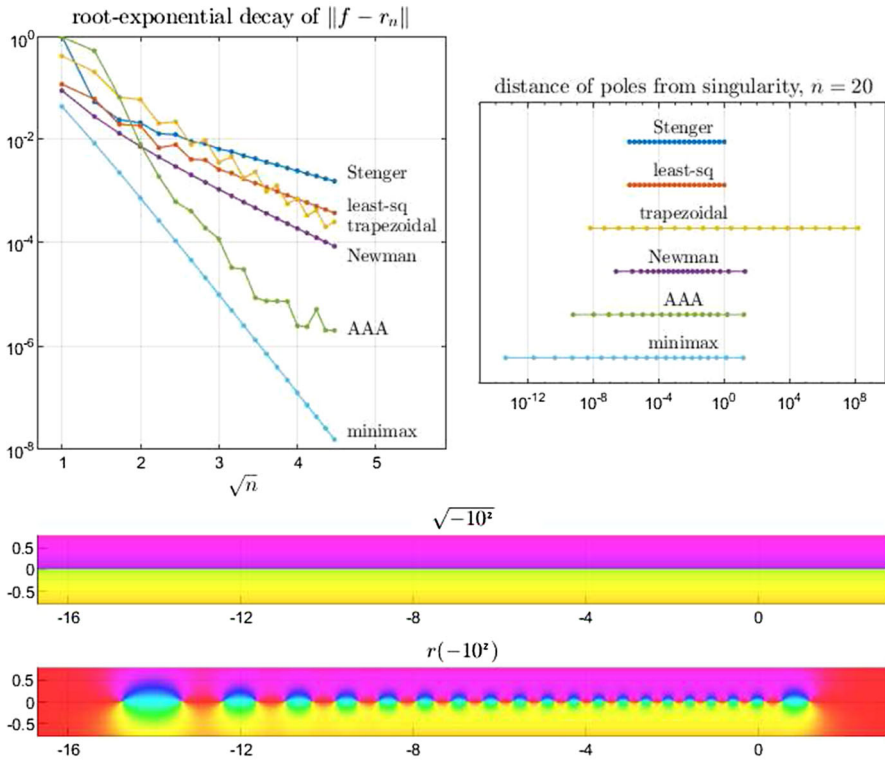
Fig. 1 Root-exponential convergence of six kinds of degree n rational approximations of $f(x) = \sqrt{x}$ on $[0, 1]$ as $n \to \infty$. On the upper-left, the asymptotically straight lines on this log scale with $\sqrt{n}$ on the horizontal axis (except for AAA) show the root-exponential effect. On the upper-right, the distances of the poles in $(-\infty, 0)$ from the singularity at $x = 0$ show the exponential clustering. Below, phase portraits in the complex plane of the square root function (the standard branch) and its degree 20 minimax approximation on $[0, 1]$, after an exponential change of variables, show how a branch cut is approximated by interlacing exponentially clustered poles and zeros. In the standard fashion for phase portraits as described in [65], red before yellow going counterclockwise indicates a zero, and yellow before red indicates a pole. We use 10^z instead of e^z to enable comparison with the axis labels in the images above

The derivation by the trapezoidal rule just given explains in a general way why root-exponential convergence is achievable for a wide range of problems with endpoint singularities. With any exponentially graded discretization, there will be errors associated with finite grid sizes and errors associated with truncation of an infinite series. If both sources of error follow an exponential dependence, then an optimal balance with step sizes scaling with $1/\sqrt{n}$ can be expected to lead to a root-exponential result. Such effects are familiar in the analysis of hp discretizations of partial differential equations when the step sizes h and orders p of multiscale discretizations are balanced to achieve optimal rates of convergence near corners [17,37]. An anonymous referee has suggested that there may be an analogy between the tapering property for rational approximants and the optimal rate of decrease of the order p in hp-finite element discretizations as a singularity is approached.

A drawback of the trapezoidal approximation is that its derivation depends on the precise spacing of the poles, since it relies on the property that the trapezoidal rule is exponentially accurate in this special case [61]. The curves labeled Stenger in Fig. 1 come from a more flexible alternative approach, also proposed by Stenger [47], where we fix n distinct poles $p_k \in (-\infty, 0)$, $1 \le k \le n$ and $n + 1$ interpolation points $x_k \in [0, 1]$, $0 \le k \le n$, and then take r to be the unique rational function of degree n with these poles that interpolates $f(x)$ in these points. The theory of rational interpolation with preassigned poles was developed by Walsh [63] and will be discussed in Sect. 4. For our problem of approximation on [0, 1] with a singularity at $x = 0$, a good choice is to take $x_0 = 0$ and $x_k = -p_k$ for $k \ge 1$. In particular, our *Stenger approximant*[1] is the rational function r resulting from the choices

$$- p_k = x_k = \exp(-(k - 1)h), \quad 1 \le k \le n, \tag{3}$$

with $h = O(1/\sqrt{n})$. Figure 1 takes $h = \pi/\sqrt{n}$.

Interpolation is important for theoretical analysis, but for practical computation, discrete least-squares fitting on a sufficiently fine grid is more robust and more accurate, since it does not require knowledge of good interpolation points. The least-squares data of Fig. 1 come from fixing the same exponentially clustered poles as in (3), but now choosing approximation coefficients by minimizing the least-squares error $f - r$ on a discretization of [0, 1] by standard methods (MATLAB backslash). As always when discretizing near singularities, we use an exponentially graded mesh (logspace(-12, 0, 2000)), and a weight function $w(x) = \sqrt{x}$ is introduced in the discrete least-squares problem so that it approximates a uniformly weighted problem on the continuum. The error curve $r(x) - \sqrt{x}$ for $x \in [0, 1]$ for this approximation (not shown) approximately equioscillates between $n + 2$ extrema, indicating that it is a reasonable approximation to the best L^∞ approximation with these fixed poles.

The minimax data in Fig. 1 correspond to the true optimal (real) approximations, rational approximations with free poles. Here the error curve equioscillates between $2n + 2$ extrema [58], and the error is approximately squared compared with the other curves; the asymptotic convergence rate is $\exp(-\pi\sqrt{2n})$ [43,62].

Computing minimax approximations, however, can be challenging [8], and on a complex domain they need not even be unique [18]. This brings us to the data in the figure for AAA (adaptive Antoulas–Anderson) approximation, a fast method of near-best rational approximation introduced in [27]. AAA approximation is at its least robust on real intervals, as reflected in the erratic data of the figure, whereas for more complicated problems and in the complex plane, it is often the most practical method for rational approximation.

This concludes our discussion of Fig. 1. The next figure, Fig. 2, illustrates that these effects are not confined to approximation on a real interval or to the function $\sqrt{x}$. The figure presents data for four further examples of minimax approximations. One set of data shows approximation of $x^{1/\pi}$ on [0, 1], with the value $1/\pi$ chosen to dispel any thought that rational exponents might be special. This problem requires

---

[1] Stenger considered rational approximations of this kind, though not in this precise setting of a finite interval with just one endpoint singularity.
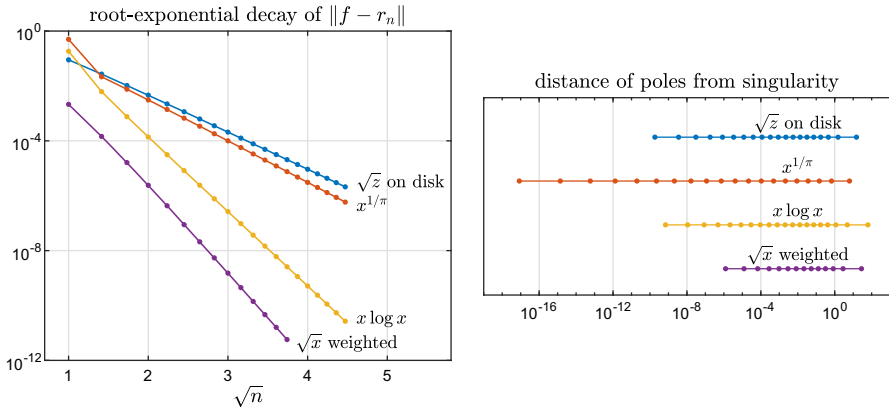
Fig. 2 Four more minimax approximations, showing the same root-exponential convergence and exponential clustering of poles as in Fig. 1. Two involve the functions $x^{1/\pi}$ and $x \log x$ on [0, 1], one involves x on [0, 1] but with the ∞-norm weighted by x, and one involves $\sqrt{z}$ on the disk about $\frac{1}{2}$ of radius $\frac{1}{2}$. In the right image, n takes its final value from the left image for each problem, 14 for the weighted approximation and 20 for the other cases
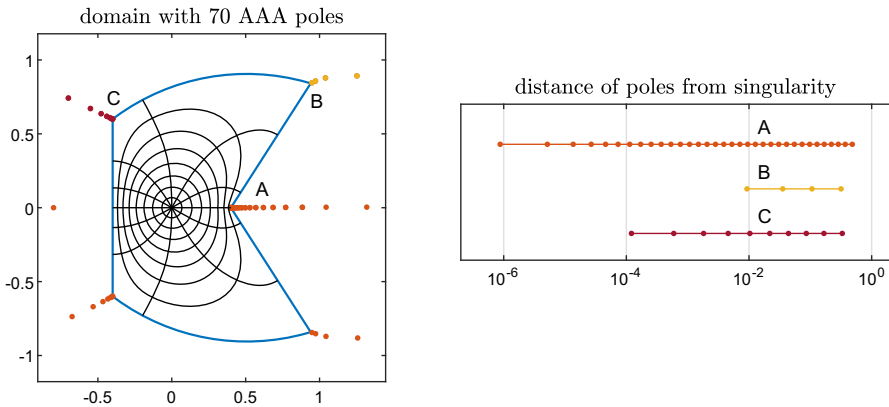

Fig. 3 The conformal map of a circular pentagon onto the unit disk has been computed and then approximated numerically by a rational function of degree 70 [14,60] by the AAA algorithm. The poles cluster exponentially at the corners, where the map is singular

poles particularly close to the singularity since the exponent is so small. Another set of data shows approximation of $x \log x$ on [0, 1]. With a much weaker singularity, this problem shows higher approximation accuracy. A third shows approximation of $\sqrt{x}$ again, but now it is weighted minimax approximation, with a weight function x (and the error measured is now the weighted error, notably smaller than before). Finally the fourth set of data shows minimax approximation of $\sqrt{z}$ on the complex disk $\{z : |z - \frac{1}{2}| < \frac{1}{2}\}$.

Figure 3 turns to our first problem of scientific computing. Following methods presented in [14,60], a region E in the complex plane bounded by three line segments and two circular arcs has been conformally mapped onto the unit disk, and the map
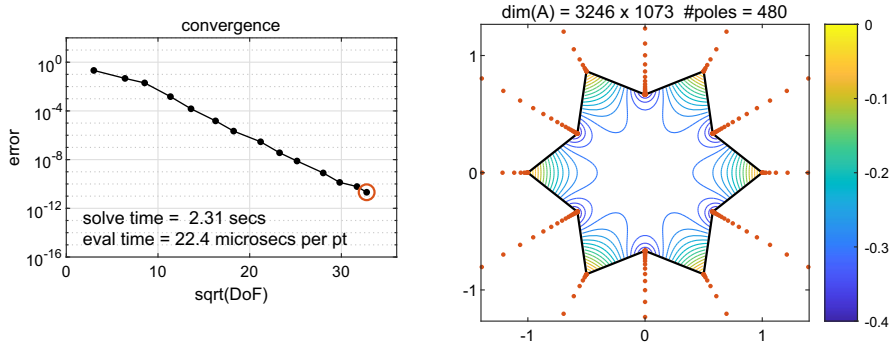
Fig. 4 Example of the lightning Laplace solver [15,16] as implemented in the code `laplace.m` [59]. For each number of degrees of freedom (DoF), poles are clustered exponentially near the 12 corners of the domain E, and the numbers are increased until a solution to 10-digit accuracy is obtained in the form of a rational function with 480 poles. This takes 2.3 s on a laptop, and subsequent evaluations take 22 μs per point, with the accuracy of each evaluation guaranteed by the maximum principle

has then been approximated to about eight digits of accuracy by AAA approximation, which finds a rational function with $n = 70$. This process is entirely adaptive, based on no a priori information about corners or singularities, yet it clusters the poles near the corners just as in Figs. 1 and 2. Many poles cluster at the strong singularity A and only a few at the weak singularity B. Note that the poles lie asymptotically on the bisectors of the external angles. This effect is well known especially from the theory of Padé approximation as worked out initially by Stahl [44,49]. Optimal approximations line up their poles along curves which balance the normal derivatives of a potential gradient on either side, and evidently the AAA method comes close enough to optimal for the same effect to appear.

We finish this section with a look at lightning solvers for PDEs in two-dimensional domains, introduced in 2019 and applied to date to Laplace [15,16,59], Helmholtz [16], and biharmonic equations (Stokes flow) [6]. In the basic case of a Laplace problem $\Delta u = 0$, the idea is to represent the solution on a domain E as $u(z) \approx \operatorname{Re} r(z)$, the real part of a rational function with no poles in E that approximates the boundary data to an accuracy typically of 6–10 digits. The rational functions have preassigned poles that cluster exponentially at the corners, where the solution will normally have singularities [22,64], and the name "lightning" alludes to this exploitation of the same mathematics that makes lightning strike objects at sharp corners. Coefficients for the solution are found by least-squares fitting, making this an approximation process of the same structure as in the least-squares example of Fig. 1. The difference is that the approximations are now applied to give values of $u(z)$ in the interior of the domain E, where it is not known a priori. See Fig. 4 for an example on a "snowflake" with boundary data $\log|z|$.

Lightning solvers have been generalized to the Helmholtz equation $\Delta u + k^2 u = 0$ [16] and the biharmonic equation $\Delta^2 u = 0$ [6], as illustrated in Fig. 5. In the Helmholtz case, poles $(z - z_k)^{-1}$ of rational functions become singularities of complex Hankel functions $H_1(k|z - z_k|) \exp(\pm i \arg(z - z_k))$, and the biharmonic case is handled by the Goursat reduction $u(z) = \operatorname{Im}(\bar{z} f(z) + g(z))$ to a coupled pair of analytic
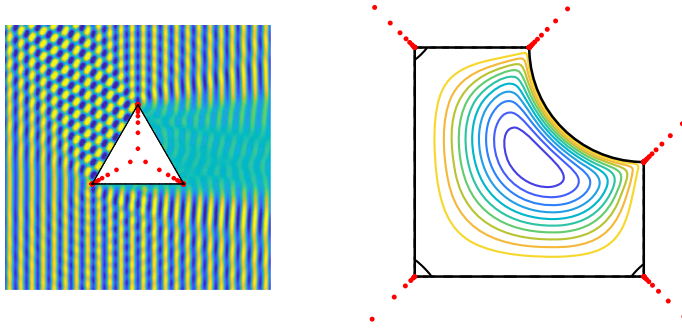
Fig. 5 Lightning solvers have been generalized to the two-dimensional Helmholtz (left) [16] and biharmonic equations (right) [6]. The Helmholtz image shows a plane wave incident from the left scattered from a sound-soft equilateral triangle. The biharmonic image shows contours of the stream function for Stokes flow in a cavity driven by a quarter-circular boundary segment rotating at speed 1 and with zero velocity on the remainder of the boundary. The black contours in the corners, representing the stream function value $\psi = 0$, delimit counter-rotating Moffatt vortices. Tapered exponentially clustered singularities are used in both computations

functions f and g, each of which is approximated by its own rational function. The mathematics of lightning methods for Helmholtz and biharmonic problems has not yet been worked out fully, and the analysis given in Sect. 5 applies just to the Laplace case.

Although it is not the purpose of this article to give details about lightning PDE solvers, they are at the heart of our motivation. Usually in approximation theory, minimax approximations are investigated as an end in themselves, and the locations of their poles may be examined as an outgrowth of this process; a magnificent example is [42]. Here, the order is reversed. Our aim is to exploit an understanding of how poles cluster to construct approximations on the fly to solve problems of scientific computing.

3 Tapered exponential clustering

In the last section, 13 plots were presented of the distances of poles to singularities on a log scale, the right-hand images of Figs. 1, 2, and 3. All showed exponential clustering, and all but three showed a further effect which we call *tapered exponential clustering*, the main subject of the rest of this paper: on the log scale, the spacing of the poles grows sparser near the singularity. This was also colorfully evident in the phase portrait at the bottom of Fig. 1. The three exceptions were the Stenger, least-squares, and trapezoidal approximations of Fig. 1, all of which are based on poles preassigned with strictly uniform exponential clustering. These examples illustrate that tapering of the pole distribution is not necessary for root-exponential convergence. A fourth set of data in Fig. 1 also involves preassigned poles, the Newman data, and some tapering is apparent in this case.

Figure 6 shows the nine remaining examples of exponential clustering of poles from Figs. 1, 2 and 3, the ones with free poles, presenting the distances $\{d_k\}$ of the poles
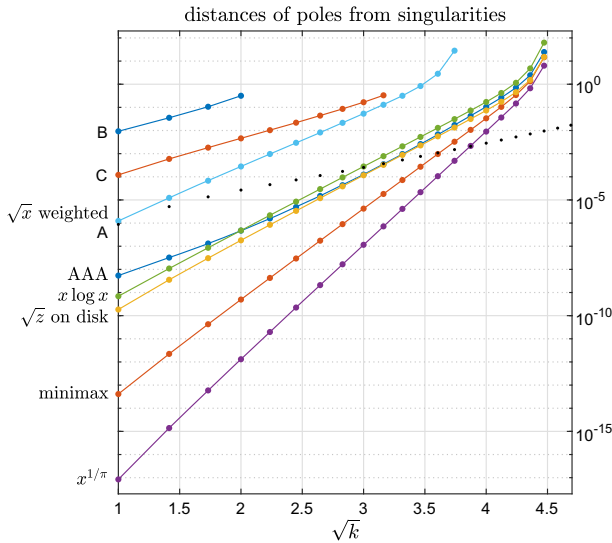
Fig. 6 Tapered exponential clustering of poles near singularities for the nine examples with free poles from Figs. 1, 2 and 3 of the last section. The crucial feature is that the curves appear straight with this horizontal axis marking $\sqrt{k}$ rather than k, where $\{d_k\}$ are the sorted distances of the poles from the singularities. The data for the poles at vertex A of Fig. 3 have been deemphasized to diminish clutter (black dots), since they lie at such a different slope from the others

from their nearest singularities on a log scale. What is immediately apparent is that all the curves look straight for smaller values of k. Note that five of them stop at $n = 20$, one at $n = 14$, and the remaining three, from the approximation of a conformal map of Fig. 3, at different values determined adaptively by the AAA algorithm.

Yet the horizontal axis in Fig. 6 is not k but $\sqrt{k}$. Plotted against k (not shown), the data would look completely different. Evidently in a wide range of rational approximations, both best and near-best, the distances $\{d_k\}$ of poles to singularities is well approximated by the formula

$$\log d_k \approx \alpha + \sigma\sqrt{k} \tag{4}$$

for some constants α and σ, that is,

$$d_k \approx \beta \exp(\sigma\sqrt{k}) \tag{5}$$

for some β and σ. It is interesting that a hint of tapering can also be seen in the optimal finite differencing grids plotted in Figure 3.3 of [20].

To make sense of the $\sqrt{k}$ scaling, let us remove the exponential from the problem by defining a distance variable $s = \log d$, thereby transplanting an interval such as $d \in [0, 1]$ to $s \in (-\infty, 0]$. We ask, what can be said of the density $\rho(s)$ of poles with respect to s? If $\rho(s)$ were constant, this would correspond to a uniform exponential distribution of poles, requiring an infinite number of poles since s goes to $-\infty$. So some kind of cutoff of $\rho(s)$ to 0 must occur as $s \to -\infty$. An abrupt cutoff, as with

UNIFORM   TAPERED

$\rho(s)$ $\rho(s)$

s s

$\int \rho(s)$ $\int \rho(s)$

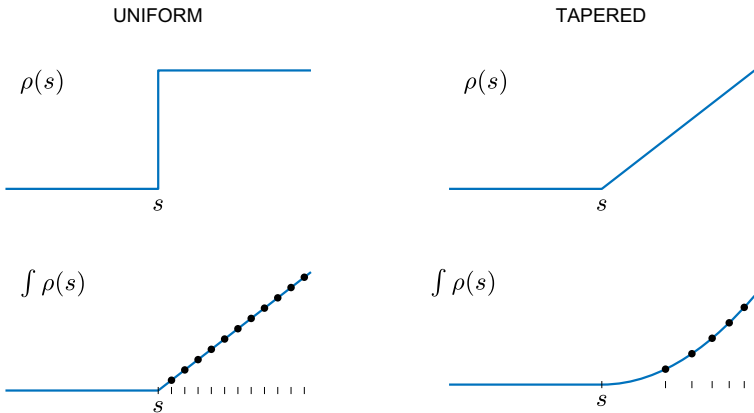s s

Fig. 7 The algebra of exponential clustering. With respect to the variable $s = \log d$, where d is the distance to the singularity, the simplest exponential clustering of poles would have uniform density $\rho(s)$ down to a certain value and then cut off abruptly (left column). A tapered distribution cuts off linearly instead (right column), resulting in poles exponentially clustered in the $\sqrt{k}$ fashion seen in Fig. 6

the Stenger, trapezoidal, and least-squares distributions of Fig. 1, leads to a linear cumulative distribution, as shown in the left column of Fig. 7. By contrast, a linear cutoff gives a quadratic cumulative distribution, as shown in the right column, and when this is inverted, the result is the $\sqrt{k}$ distribution we have observed.

Thus the straight lines of Fig. 6 can be explained if pole density functions $\rho(s)$ for good rational approximations tend to take the form sketched in the upper-right of Fig. 7. (Aficionados of deep learning may call this the "ReLU" shape.) In Sect. 5 we will explain why this is the case and continue the story of Fig. 7 in Fig. 11.

We have not presented data in this section for lightning PDE solutions, but it was in this context that we first became aware of the importance of tapered exponential clustering. In the course of the work leading to [15], the first author noticed that although straight exponential spacing of preassigned poles gave root-exponential convergence, better efficiency could be achieved if the resulting approximations were re-approximated a second time by the AAA algorithm. On examination it was found that the AAA approximations had poles in a tapered distribution, just like cases A–C of Fig. 6. The model (4)–(5) was developed empirically in this context, with $\sigma \approx 4$ found to be an effective choice. This became the formula for preassignment of poles in the lightning Laplace software [59], where it improved the overall speed by a good factor, and it appears as equation (3.6) in [15].

4 Hermite integral formula and potential theory

The basic tool for estimating accuracy of rational approximations is the Hermite integral formula [23,63]. In this section we review how this formula leads to the use of potential theory [32], and in particular the quantity known as the condenser capacity, for approximations of analytic functions. Building on the work of Walsh [63], these

ideas began to be developed by Gonchar and Rakhmanov in the Soviet Union not long after the appearance of Newman's paper [12,13].

The following statement is adapted from Thm. 8.2 of [63].

Theorem 1 *Let Ω be a simply connected domain in $\mathbf{C}$ bounded by a closed curve Γ, and let f be analytic in Ω and extend continuously to the boundary. Let distinct interpolation points $x_0, \ldots, x_n \in \Omega$ and poles $p_1, \ldots, p_n$ anywhere in the complex plane be given. Let r be the unique degree n rational function with simple poles at $\{p_k\}$ that interpolates f at $\{x_k\}$. Then for any $x \in \Omega$,*

$$f(x) - r(x) = \frac{1}{2\pi i} \int_\Gamma \frac{\phi(x)}{\phi(t)} \frac{f(t)}{t - x} dt, \tag{6}$$

where

$$\phi(z) = \prod_{k=0}^{n}(z - x_k) \Big/ \prod_{k=1}^{n}(z - p_k). \tag{7}$$

To see how this theorem is applied, let Ω be a simply connected domain bounded by a closed curve Γ, as indicated in Fig. 8 (see also Fig. 9 in the next section), and let f be analytic in Ω and extend continuously to Γ. Suppose f is to be approximated on a compact set $E \subset \Omega \cup \Gamma$, which in this section we take to be disjoint from Γ. Theorem 1 implies that for any $x \in E$,

$$|f(x) - r(x)| \leq C\tau, \tag{8}$$

where C is a constant independent of n and τ is the ratio

$$\tau = \frac{\max_{z \in E} |\phi(z)|}{\min_{z \in \Gamma} |\phi(z)|}. \tag{9}$$

If ϕ is much smaller on E than on Γ, then τ and hence $f - r$ must be small.

Figure 8 gives an idea of how this can happen. In each image, the red dots on Γ represent a good choice of poles $\{p_k\}$ and the blue dots on the boundary of E a corresponding good choice of interpolation points $\{x_k\}$. Consider first the upper-left image, where E and Γ define a circular annulus. The equispaced configurations of $\{p_k\}$ and $\{x_k\}$ ensure that τ will decrease exponentially as $n \to \infty$. To see this, in view of (7), we define

$$u(z) = n^{-1} \sum_{k=0}^{n} \log |z - x_k| - n^{-1} \sum_{k=1}^{n} \log |z - p_k|. \tag{10}$$

This is the potential function generated by $n + 1$ negative point charges of strength n^{-1} at the interpolation points and n positive point charges of strength $-n^{-1}$ at the poles. Then $\exp(nu(z)) = |\phi(z)|$, and therefore

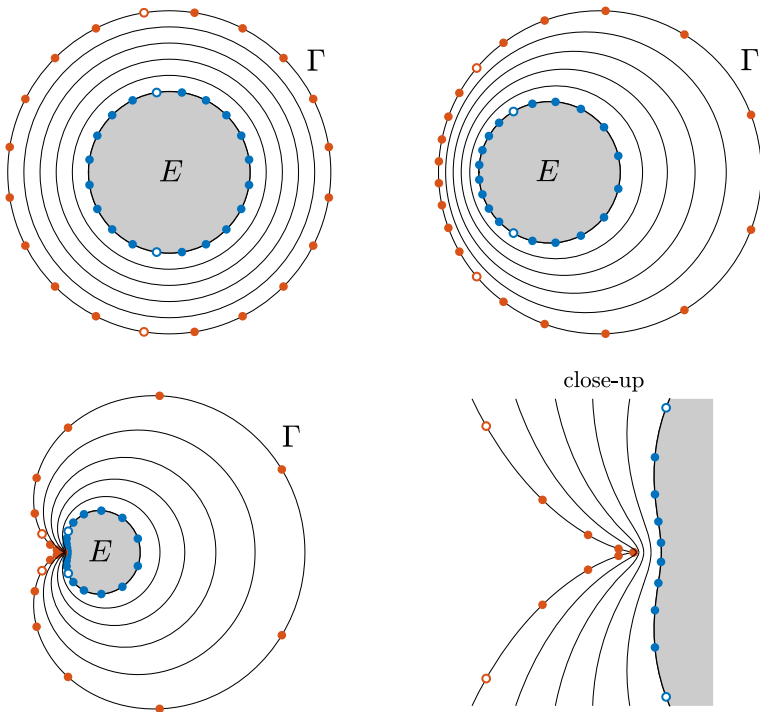$$\tau = \exp\left(-n\left[\min_{z \in \Gamma} u(z) - \max_{z \in E} u(z)\right]\right). \tag{11}$$

Fig. 8 Potential theory and rational approximation. In each image, the shaded region is an approximation domain E for a function f analytic in the region Ω bounded by Γ. If we think of the poles of an approximation $r \approx f$ as positive point charges and the interpolation points as negative point charges, then a minimal-energy equilibrium distribution of the charges gives a favorable configuration for approximation. This is a discrete problem of potential theory that becomes continuous in the limit $n \to \infty$, enabling one to take advantage of invariance under conformal maps. In these images E and Γ are disjoint and the convergence is exponential, but the third domain and its close-up illustrate the clustering effect, which will become more pronounced as the gap shrinks to zero. The pairs of interpolation points and poles marked by hollow dots delimit one half of the total, highlighting how both sets of points accumulate close to the singularity

For τ to be small, we want u to be uniformly bigger on Γ than on E. Finding the best such configuration is an extremal problem that will be approximately solved if the points are placed in an energy-minimizing equilibrium position. In each of the images of Fig. 8, the points are close to such an equilibrium. Each charge is attracted to the charges of the other color, but repelled by charges of its own color.

Finding an optimal configuration (for the given choice of Γ) is complicated for finite n, but the problem becomes cleaner in the limit $n \to \infty$, and this is where the power of potential theory is fully revealed. We now imagine continua of interpolation points and poles defined by a signed measure μ supported on E, where it is nonpositive with total mass -1, and on Γ, where it is nonnegative with total mass 1. (We assume that E is well-enough behaved that complications such as components of zero capacity do not come into play.) It can be shown that there is a unique measure of this kind that

minimizes the energy

$$I(\mu) = -\int \int \log |z - t| d\mu(z) d\mu(t), \qquad (12)$$

with associated potential function

$$u(z) = -\int \log |z - t| d\mu(t), \qquad (13)$$

and u takes constant values $u_E < 0$ on E and $u_\Gamma = 0$ on Γ. The minimum $I_{\min} = \inf_\mu I(\mu)$ is known to be positive, and for minimax degree n rational approximations r_n^* one has exponential convergence as $n \to \infty$ at a corresponding rate:

$$\limsup \| f - r_n^* \|^{1/n} \leq \exp(-I_{\min}). \qquad (14)$$

(The actual rate is in fact twice as fast as this, $\exp(-2I_{\min})$, for functions whose singularities in the complex plane are just isolated algebraic branch points [23, p. 93], [39].)

The reciprocal of $I_{\min}$ is known as the *condenser capacity* for the (E, Γ) pair, a term that reflects an electrostatic interpretation of the approximation problem. In electronics, capacitance is the ratio of charge to voltage difference. A capacitor has high capacitance if its positive and negative plates are close to one another, so that the attraction of charges of opposite sign enables a great deal of charge to be accumulated on them without the need for much of a voltage difference. For fast-converging rational approximation, on the other hand, we want Γ and E to be far apart, corresponding to a *small* amount of charge relative to the voltage difference, hence small numbers of poles and interpolation points needed to achieve a given ratio τ.

We can now see how the second and third images of Fig. 8 were drawn. They were obtained by conformal transplantation, exploiting the invariance of problems of potential theory under conformal maps. The eccentric domain of the second image comes from a Möbius transformation, and the pinched domain of the third image comes from a further squaring. The blue and red points obtained as conformal images of equispaced points in the symmetric annulus are known as Fejér–Walsh points [45].

One might wonder, for arguments of this kind, is it necessary to place the poles of r on the boundary of the region of analyticity of f? In fact, Γ does not have to lie as far out as that boundary, nor do the poles have to be on Γ, for as stated in Theorem 1, the integral representation (6) is valid for any placement of the poles. Asymptotically as $n \to \infty$, however, it is known that the convergence rate cannot be improved by placing poles beyond the region of analyticity of f [23]. A special choice is to put all the poles at $x = \infty$, in which case rational approximation reduces to polynomial approximation, still with exponential convergence though at a lower rate than in (14).
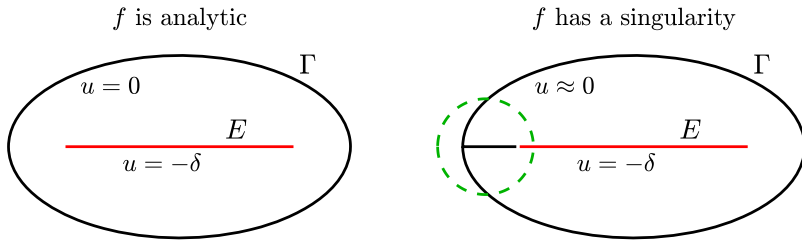
Fig. 9 Two kinds of problems of rational approximation of a function f on a domain E. On the left (Sect. 4), f is analytic on E and poles can be placed on a contour Γ enclosing E in the region of analyticity: convergence is exponential with accuracy on the order of $\exp(-n\delta)$ for a constant $\delta > 0$. On the right (Sect. 5), f has a singularity at a point z_c on the boundary of E, and Γ must touch E at z_c: convergence is root-exponential with accuracy of order $\exp(-n\delta)$ again, but now with δ diminishing at the rate $1/\sqrt{n}$ as $n \to \infty$. In the circled region, the potential makes the transition from $u_\Gamma \approx 0$ to $u_E = -\delta$

5 Explanation of tapered exponential clustering

Now we examine how the analysis of the last section must change for approximations with singularities. There is a considerable specialist literature here by authors including Aptekarev, Saff, Stahl, Suetin, and Totik [2,23,34,36,43,44,49], which investigates certain best approximations in detail. Our emphasis is on the broad ideas applicable to near-best approximations too.

From Fig. 8 it is clear that potential theory should give some insight when f has a singularity on the boundary of E. The lower pair of images shows clustering of poles where Γ has a cusp close to the boundary of E, and as the cusp is brought closer to E, the clustering will grow more pronounced. However, the argument we have presented breaks down when Γ actually touches E. The situation is sketched in Fig. 9. Physically, this would be a capacitor of infinite capacitance, implying that an equipotential distribution u with a nonzero voltage difference would require an infinite quantity of charge. Mathematically, the estimate (8) fails because τ cannot be smaller than 1.

To see what happens in such cases, we can examine the function ϕ computed numerically for an example problem. The left column of Fig. 10 shows error curves in type (9, 10) minimax approximation of $\sqrt{x}$ on [0.01, 1] (above) and [0, 1] (below). (Type (m, n) means numerator degree at most m and denominator degree at most n; we choose these parameters rather than (n, n) to make the plots slightly cleaner.) The curves each equioscillate between $m + n + 2 = 21$ extrema, and in the lower curve, on the semilogx scale, we see the wavelength increasing as $x \to 0$. As a minimax approximation with free poles, this rational function has $m + n + 1 = 20$ points of interpolation rather than the usual number $n = 10$ for an approximation with preassigned poles.

The right column of Fig. 10 shows the function $|\phi|$ plotted on the approximation interval (the lower blue curve) and on the important portion $[-1, 0]$ of the integration contour Γ (the upper red curve). More precisely, for these plots we have taken the numerator of (7) to range over just the $n + 1$ interpolation points marked by red dots in the left column, this being the usual number for near-best but not for exact best
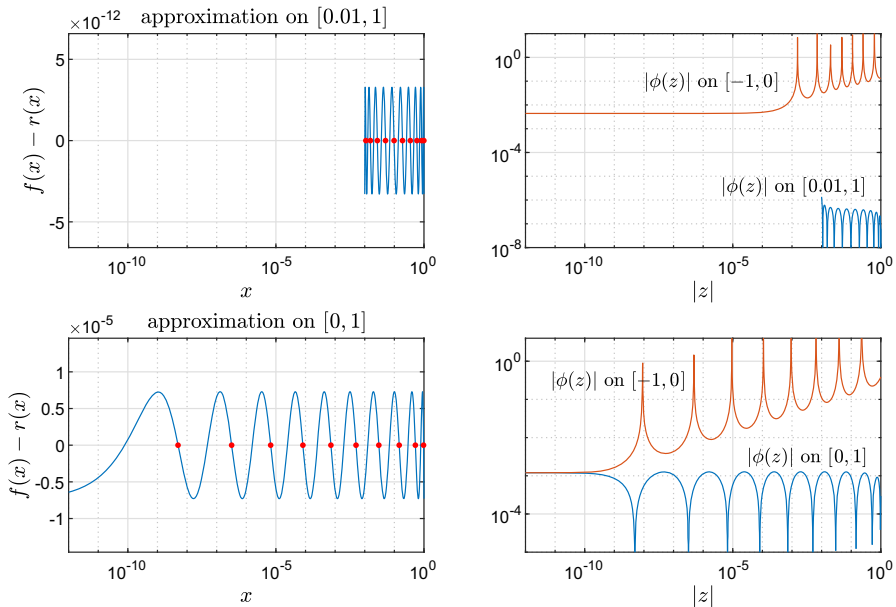
Fig. 10 On the left, error curves in type $(9, 10)$ minimax approximation of $\sqrt{x}$ on $[0.01, 1]$ and $[0, 1]$. On the right, plots of $\phi(z)$ as defined by (7) on these approximation intervals and on $[-1, 0]$. The curves in the upper-right image show a reasonable approximation to constant values on $[-1, 0]$ (upper curve) and on $[0.01, 1]$ (lower curve), but in the lower-right image, nothing like constant behavior of $|\phi(z)|$ on $[-1, 0]$ is evident. We explain this by noting that what matters to the accuracy of an approximation is the integral (15) of $|\phi(x)/\phi(t)|$ with respect to $t \in \Gamma$, not its maximum. Taking advantage of this property, poles and interpolation points distribute themselves more sparsely near the singularity, freeing more of them to contribute to the approximation further away—the phenomenon of tapered exponential clustering

approximation. Including all $2n + 1$ points would introduce a further slope in both the blue and red curves that would distract from the main point. In the upper image of the right column, for $[0.01, 1]$, the curves reveal a reasonable approximation to what the last section has led us to expect from potential theory. The envelope of the blue curve has approximately constant magnitude, and this is about five orders of magnitude below the envelope of the red curve, also of approximately constant magnitude. Thus the ratio τ of (9) is far below 1, and the estimate (8) serves to bound the approximation error. (The actual error is about the square of this bound since we have omitted half the interpolation points.)

The lower image, which is a centerpiece of this paper, tells a strikingly different story. Here again the envelope of the blue curve is flat, showing the invariance with respect to x we expect in a minimax approximation. The envelope of the red curve for $|\phi(z)|$ on $[-1, 0]$, however, is now tilted at an angle on these log-log axes, showing a steady closing of the gap between the curves as z moves from -1 to 0. Clearly in this case $[-1, 0]$ is not at all an interval where the envelope of $|\phi|$ is constant.

To understand the linearly closing gap in Fig. 10, we note that what fails in the analysis of the last section for an approximation problem with a singularity is not the

Hermite integral,

$$f(x) - r(x) = \frac{1}{2\pi i} \int_\Gamma \frac{\phi(x)}{\phi(t)} \frac{f(t)}{t - x} dt, \tag{15}$$

but the estimate (8) we derived from it. Implicitly (8) came from bounding the right-hand side of (15) by Hölder's inequality,

$$|f(x) - r(x)| \le \frac{1}{2\pi} \left\| \frac{\phi(x)}{\phi(t)} \right\|_\infty \left\| \frac{f(t)}{t - x} \right\|_\infty \|1\|_1, \tag{16}$$

where the ∞- and 1- norms are defined over $t \in \Gamma$. (The norm $\|1\|_1$ is equal to the length of Γ.) When Γ and E are disjoint, the first ∞-norm in (16) is exponentially small as $n \to \infty$ and the second is bounded uniformly in x. However, these properties fail as Γ and E touch. We can rescue the argument by noting that $|\phi(x)/\phi(t)|$ does not have to be small for all t so long as its integral is small. More precisely, the quantity $f(t)/(t - x)$ of (16) may not be bounded as $t, x \to z_c$ but $f(t)|t - z_c|^{1-\alpha}/(t - x)$ will be bounded if we assume $f(t - z_c) = O(|t - z_c|^\alpha)$ for some constant $\alpha \in (0, 1]$. So what actually matters is that the integral of $|t - z_c|^{\alpha-1}|\phi(x)/\phi(t)|$ should be small, and we accordingly replace (16) by the alternative Hölder estimate

$$|f(x) - r(x)| \le \frac{1}{2\pi} \left\| \frac{\phi(x)}{\phi(t)} |t - z_c|^{\alpha-1} \right\|_1 \left\| \frac{f(t)}{t - x} |t - z_c|^{1-\alpha} \right\|_\infty. \tag{17}$$

For the simplest model problem let us assume that the singularity lies at $z_c = 0$, with E to the left along the negative real axis and the part of the contour Γ that matters being $[0, 1]$, and let the domain be scaled so that the envelope of $|\phi(x)|$ is approximately 1 for $x \in E$. We want $|\phi(t)|$ to be at most 1 for $t < 0$ and at least $(t/\varepsilon)^\alpha$ for $t > \varepsilon$, where ε is the distance of the closest pole to the singularity. See the upper-right image of Fig. 11. Immediately below that image is a sketch of the model problem suggested by defining $u(t) = n^{-1} \log \phi(t)$ as in (10): find a harmonic function $u(t)$ in the upper half t-plane such that

$$u(t) = \begin{cases} 0, & t \le \varepsilon, \\ \alpha n^{-1} \log(t/\varepsilon), & t > \varepsilon. \end{cases} \tag{18}$$

We now make the change of variables $s = \log t$, which transplants the Laplace problem to the infinite strip $S = \{s \in \mathbf{C} : 0 < \mathrm{Im}\, s < \pi\}$, as sketched in the (3, 2) position of the figure: find a harmonic function $u(s)$ in S satisfying

$$u(s) = \begin{cases} 0, & \mathrm{Im}\, s = \pi, \\ 0, & \mathrm{Im}\, s = 0 \text{ and } \mathrm{Re}\, s \le \log \varepsilon, \\ \alpha n^{-1}(s - \log \varepsilon), & \mathrm{Im}\, s = 0 \text{ and } \mathrm{Re}\, s > \log \varepsilon. \end{cases} \tag{19}$$

(To keep matters simple we have slightly abused notation by retaining the same variable name u.) This change of variables is convenient mathematically, and it is also important conceptually, since it is well known that influences on harmonic functions often decay exponentially with distance along a strip (for biharmonic functions this is called St.
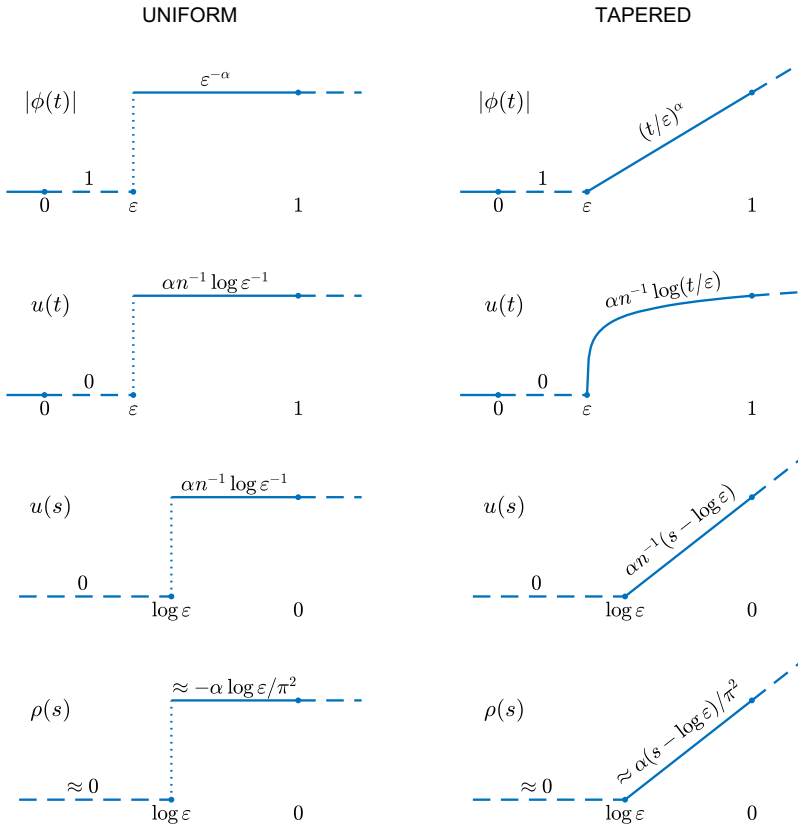
UNIFORM   TAPERED

Fig. 11 The potential theory of exponential clustering, in continuation of Fig. 7. (We comment on the right column; the left column is analogous.) The first two rows show the function $|\phi(t)|$ of (7) and the associated potential $u(t) = n^{-1} \log |\phi(t)|$ for the model problem (18). The third row shows the behavior along the real s-axis after the change of variables to $s = \log t$ and the model problem (19); the domain is now the infinite strip $0 < \operatorname{Im} s < \pi$, with $u = 0$ for $\operatorname{Im} s = \pi$. The final row shows the charge density $\rho(s) = (n/\pi)(\partial/\partial n)u(s)$, where $\partial/\partial n$ is the external normal derivative of u on the boundary of the strip. The intervals that matter (emphasized by solid rather than dashed lines) are $\varepsilon < |t| < 1$ in the t variable and $\log \varepsilon < \operatorname{Re} s < 0$ in the s variable. Smaller values of $|t|$ and s contribute negligibly to the integral (15), and larger values are far from the singularity

Venant's principle). Consequently, if ε is small, the solution to a Laplace problem for $\log \varepsilon \ll \operatorname{Re} s \ll 0$ will be essentially (though not exactly) determined by the boundary conditions in that region. This just matches what we need for the model problem as posed in the original t variable, where behavior for $|t|$ of order ε or less is unimportant because it contributes negligibly to the integral (15) and behavior for $|t|$ of order 1 or more is unimportant because it is far from the singularity under investigation.

So we address our attention to (19). An exact solution can be obtained via the Poisson integral formula for an infinite strip [66],

$$u(x + iy) = \frac{\alpha n^{-1}}{2\pi} \int_0^\infty \frac{\xi \sin(y)}{\cosh(\xi - (x - \log \varepsilon)) - \cos(y)} d\xi, \qquad (20)$$

where we have set $s = x + iy$ with $x, y \in \mathbf{R}$. (This use of x is unrelated to $x \in E$ above.) However, we do not need exactly this since the region where our model applies is $\log \varepsilon \ll \operatorname{Re} s \ll 0$. In this region, the bilinear harmonic function

$$u(x + iy) = \alpha n^{-1} \left(1 - \frac{y}{\pi}\right)(x - \log \varepsilon) \tag{21}$$

satisfies the boundary conditions and is accordingly a good approximation to the solution to (19). To determine the corresponding pole density we note that the harmonic function u has a single-layer integral representation over the boundary S

$$u(s) = \int_S h(\eta) \frac{\log |s - \eta|}{2\pi} |d\eta|$$

for some density function h, and the normal derivative of u jumps by $h(s)$ as s crosses the boundary [9, Cor. 3.29]. For our problem this jump condition becomes

$$h(s) = -2 \frac{\partial}{\partial y} u(s)$$

because of symmetry across the boundary (seen most readily in the t variable, where a configuration of poles and interpolation points on the real axis generates a function $\phi(t)$ satisfying $|\phi(\bar{t})| = |\phi(t)|$). To relate this to the pole density (imagining a continuum of poles in the limit $n \to \infty$) we write $\rho(s) = (n/2\pi)h(s)$, since the formula (10) defining u contains logarithms divided by n rather than by 2π. The result for (21) is

$$\rho(s) = -\frac{n}{\pi} \frac{\partial}{\partial y} u(x + iy) = \frac{\alpha}{\pi^2}(x - \log \varepsilon). \tag{22}$$

This linear growth as a function of x, sketched in the bottom-right image of Fig. 11, is just what we set out to explain in Fig. 7.

Let us now look at the quantitative implications of this argument, comparing uniform exponential clustering (left column of Fig. 11) with tapered exponential clustering (right column). According to our model, the integral of $\rho(s)$ over $[\log \varepsilon, 0]$ as sketched in the bottom row of images should be equal to n, the total number of poles, and this enables us to determine ε. For uniform clustering the integral is $\alpha \log^2(\varepsilon)/\pi^2$, leading to the estimates

$$\text{Closest pole: } \varepsilon \approx \exp(-\pi\sqrt{n/\alpha}), \quad \text{Accuracy: } \varepsilon^\alpha \approx \exp(-\pi\sqrt{\alpha n}). \tag{23}$$

For tapered clustering the integral is $\frac{1}{2}\alpha \log^2(\varepsilon)/\pi^2$, leading to the estimates

$$\text{Closest pole: } \varepsilon \approx \exp(-\pi\sqrt{2n/\alpha}), \quad \text{Accuracy: } \varepsilon^\alpha \approx \exp(-\pi\sqrt{2\alpha n}). \tag{24}$$

Thus, as mentioned in the abstract, our model leads to the prediction of a factor of 2 speedup with tapered clustering. It would be interesting to investigate whether,
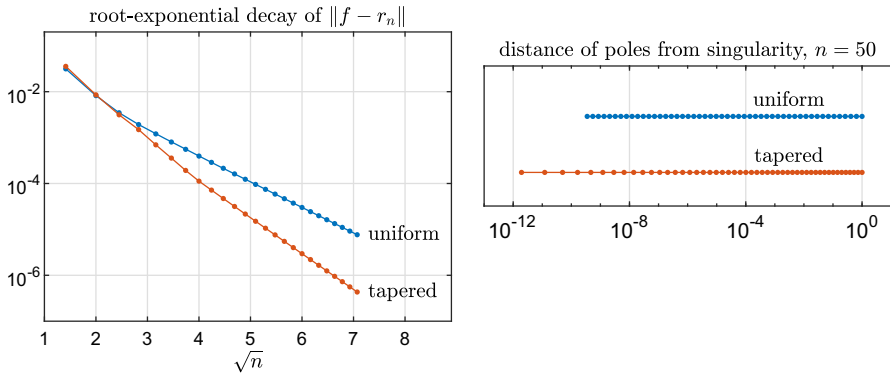
Fig. 12 Linear-minimax approximation of $f(x) = \sqrt{x}$ on $[0, 1]$ with preassigned exponentially clustered poles in $[-1, 0]$, $n = 2, 4, \ldots, 50$. Tapering the distribution makes the convergence rate approximately double, as predicted by the model of Sect. 5

for certain problems, exactly this ratio could be established theoretically in the limit $n \to \infty$.

As an example of a problem in which we may make such a comparison numerically, consider Fig. 12. These data show ∞-norm errors for rational linear-minimax approximations of $\sqrt{x}$ on $[0, 1]$ of even degrees n from 2 to 50 with preassigned exponentially clustered poles. That is, the approximations are optimal in the ∞-norm among rational functions in R_n with simple poles at the prescribed points; they are characterized by error curves equioscillating between $n + 2$ extrema. The upper curves correspond to uniformly clustered poles $p_k = -\exp(-\pi k/\sqrt{n})$, $0 \le k \le n - 1$, and the lower curves to tapered poles $p_k = -\exp(\sqrt{2}\pi(\sqrt{k} - \sqrt{n}))$, $1 \le k \le n$. The asymptotic errors appear to be about $\exp(-\sqrt{2.3n})$ for uniform clustering and $\exp(-\sqrt{4.7n})$ for tapered clustering. With $\alpha = 1/2$ for $f(x) = \sqrt{x}$, the corresponding estimates (23) and (24) are $\exp(-\sqrt{2.2n})$ and $\exp(-\sqrt{4.4n})$.

Analyses related to the argument we have presented were published by Stahl for rational minimax approximation of $|x|$ on $[-1, 1]$ and x^α on $[0, 1]$ [40–43]. For x^α Stahl gives a result that can be loosely summarized as

$$\text{Accuracy: } \varepsilon^\alpha \approx \exp(-\pi\sqrt{4\alpha n}). \tag{25}$$

(His theorem gives precise asymptotic behavior for the limit $n \to \infty$, assuming α is not an integer.) This is just what one would expect based on (24), since, as mentioned earlier, the effective value of n is doubled in the case of true minimax approximants [31]. Stahl worked essentially in the variable t rather than s, so his boundary conditions involved logarithms, as in the second image of the right column of Fig. 11. Whenever one has a Laplace problem with Dirichlet boundary data, one can interpret it as the problem of finding an equipotential distribution in the presence of an external field defined by that boundary data, and this interpretation has been carried far in approximation theory [36]. From this point of view one can say that tapered exponential clustering results from poles and zeros being slightly pushed away from a singular point by a logarithmic potential field.

6 Exponential and double exponential quadrature

In this final section we turn to another problem where exponential clustering appears. Let f be a continuous function on $[-1, 1]$. We wish to approximate the integral of f over $[-1, 1]$ by a linear combination

$$I_n = \sum_{k=1}^{n} w_k f(x_k), \tag{26}$$

where $\{x_k\}$ are distinct nodes in $[-1, 1]$ and $\{w_k\}$ are corresponding weights, in such a way that the accuracy is good even if f has branch point singularities at the endpoints. To this end, we introduce a change of variables $x = g(s)$ from the real line to $[-1, 1]$, so that the integral becomes

$$I = \int_{-1}^{1} f(x)dx = \int_{-\infty}^{\infty} f(g(s))g'(s)ds, \tag{27}$$

and we apply the equispaced trapezoidal rule. This involves an infinity of sample points in principle, but if $g'(s)$ decays rapidly, we may truncate these to an n-point rule like (2):

$$I_n = h \sum_{k=-(n-1)/2}^{(n-1)/2} f(g(kh))g'(kh). \tag{28}$$

Quadrature formulas of this kind were introduced around 1970 by Mori, Takahasi, and other Japanese researchers and also in the analysis of sinc methods by Stenger. See [19,21,46,48,52,61], as well as [24] for the history as told by Mori himself. The standard "exponential" choice of g is

$$g(s) = \tanh(s), \quad g'(s) = \operatorname{sech}^2(s), \tag{29}$$

with which (28) becomes the *tanh formula*. As in Sect. 2, we estimate the truncation error as of order $\exp(-nh)$ and the discretization error of order $\exp(-\pi^2/h)$. (The latter could be worse if f has additional singularities near $(-1, 1)$.) This gives a balance $h \approx \pi/\sqrt{n}$, with convergence rate of order $\exp(-\pi\sqrt{n})$. An estimate of this form is valid for any Hölder continuous branch point singularity; see [46, Thm. 3.4], [54, Thm. 2.1], and [61, Thm. 14.1].

Root-exponential convergence! This is much better than any algebraic order, but for practical applications on one-dimensional domains, methods of this kind often seem very wasteful, with almost all the points being used up in resolving singularities (100% of them, in the limit $n \to \infty$) [33]. A year or two after the first exponential formulas appeared, it was realized that one can do better with "double exponential" formulas. We focus on the *tanh-sinh formula* proposed by Takahasi and Mori in [53] and subsequently used and analyzed by many others including Okayama, Sugihara,
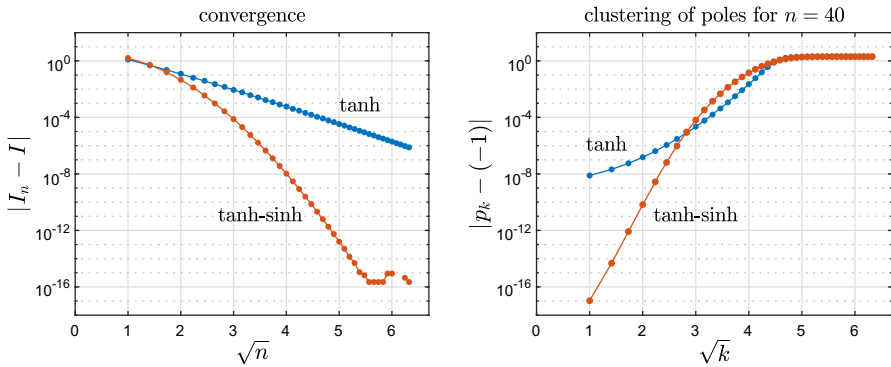
Fig. 13 On the left, root-exponential convergence of the tanh quadrature formula applied to integration of $\sqrt{1+x}$ (note the $\sqrt{n}$ axis as usual); the tanh-sinh formula converges much faster down to machine precision. On the right, the distances of nodes from poles (with a $\sqrt{k}$ axis) show uniform exponential clustering for the tanh formula with $n = 40$ and tapered exponential clustering for tanh-sinh

and Tanaka as well as Bailey and Borwein [1,25,30,50,54]. Here (29) is replaced by

$$g(s) = \tanh\left(\tfrac{\pi}{2}\sinh(s)\right), \quad g'(s) = \tfrac{\pi}{2}\cosh(s)\mathrm{sech}^2\left(\tfrac{\pi}{2}\sinh(s)\right). \tag{30}$$

Under suitable assumptions we can now estimate the truncation and discretization errors as of orders $\exp(-(\pi/2)\exp(nh/2))$ and $\exp(-\pi^2/h)$. The first of these estimates is the big improvement, for this quantity can be almost-exponentially small with a much smaller value of h than before, of order $\log(n)/n$ rather than $1/\sqrt{n}$. By *almost-exponential,* we mean of order $\exp(-Cn/\log n)$ for some $C > 0$. With this reduced value of h, the second estimate becomes almost-exponentially small too.

Figure 13 shows data for the tanh and tanh-sinh formulas. (We used the empirical choices $h = \pi/\sqrt{n}$ and $h = 1.2\log(2\pi n)/n$, respectively.) The left image plots $|I_n - I|$ against $\sqrt{n}$ for n from 1 to 40 for the integration of $f(x) = \sqrt{1+x}$. The tanh curve appears straight, confirming the root-exponential convergence, and the tanh-sinh curve bends downward, confirming that its rate is faster. The unexpected image is on the right, a plot of distances of the nodes from the endpoint $x = -1$. For tanh quadrature, these distances are uniformly exponentially spaced, appearing as a parabola on these axes. The curve for tanh-sinh quadrature, however, is almost perfectly straight. It would seem that tanh-sinh quadrature exploits tapered exponential clustering! It surprised us when we first saw curves like this. Why is there a resemblance between the tanh and tanh-sinh quadrature formulas and the phenomena of rational approximation discussed in the earlier sections of this article?

Some steps toward an answer come from a beautiful connection introduced by Gauss and exploited by Takahasi and Mori [10,51,56]: every quadrature formula can be associated with a rational approximation. Suppose first that f can be analytically continued to a neighborhood Ω of $[-1, 1]$. Then the integral can be written

$$I = \int_{-1}^{1} f(x)dx = \frac{1}{2\pi i}\int_{\Gamma} f(t)\varphi(t)dt, \tag{31}$$

where Γ is a contour enclosing $[-1, 1]$ and disjoint from it and the *characteristic function φ* is defined by

$$\varphi(t) = \int_{-1}^{1} \frac{dx}{t - x} = \log \frac{t + 1}{t - 1}. \tag{32}$$

On the other hand the quadrature sum (26) can be written

$$I_n = \frac{1}{2\pi i} \int_{\Gamma} f(t) r(t) dt, \tag{33}$$

where r is the degree n rational function defined by

$$r(t) = \sum_{k=1}^{n} \frac{w_k}{t - x_k}. \tag{34}$$

Subtracting (33) from (31) gives what we call the *Gauss–Takahasi–Mori (GTM) contour integral,*

$$I - I_n = \frac{1}{2\pi i} \int_{\Gamma} f(t)(\varphi(t) - r(t)) dt \tag{35}$$

and the corresponding error bound

$$|I - I_n| \leq \frac{1}{2\pi} \|f\|_{\infty} \|\varphi - r\|_{\infty} \|1\|_1, \tag{36}$$

which we have written in the style of (16), with the norms defined over Γ.

Equations (35) and (36) relate accuracy of a quadrature formula to an approximation problem: if the nodes and weights are such that $\varphi - r$ is small on the boundary Γ of a region where f is analytic, then $|I - I_n|$ must be small. This reasoning was applied by Takahasi and Mori to a range of quadrature formulas [51]. The function φ is analytic in the extended complex plane minus the segment $[-1, 1]$. It follows that since Γ is disjoint from $[-1, 1]$, there exist rational approximations to φ that converge exponentially on Γ as $n \to \infty$. In particular, this holds for the rational functions associated with Gauss and Clenshaw–Curtis quadrature [56], where it is convenient to take Γ in the form of an ellipse about $[-1, 1]$ with foci ± 1. It follows that both these quadrature formulas converge exponentially as $n \to \infty$ for analytic integrands (cf. [58, Thm. 19.3]).

But what if f has endpoint singularities? Now Γ must be taken to be a contour that touches $[-1, 1]$ at the endpoints, and (36) fails just as (16) did in such a case. In fact, this failure is more severe, since $\|\varphi - r\|_{\infty} = \infty$ for any r because of the logarithmic singularities of φ. The last section, however, suggests a solution. Instead of (36), we can derive from (35) the bound

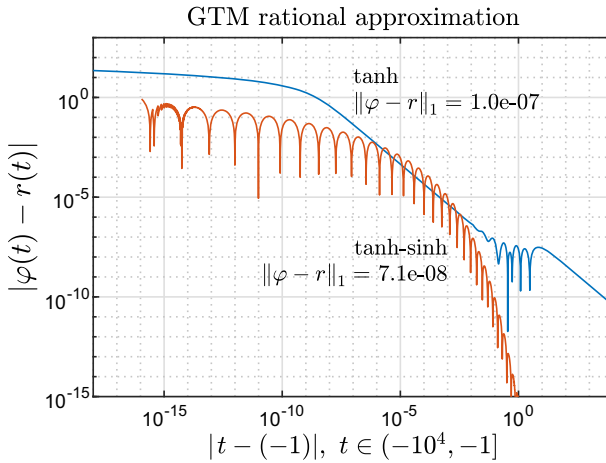$$|I - I_n| \leq \frac{1}{2\pi} \|f\|_{\infty} \|\varphi - r\|_1. \tag{37}$$

GTM rational approximation

Fig. 14 Error $|\varphi(t) - r(t)|$ as a function of distance to the left from $t = -1$ for the tanh and tanh-sinh approximations with $n = 40$ with $h = \pi/\sqrt{n}$ and $h = 1.2 \log(2\pi n)/n$, respectively. By symmetry, the same behavior would appear to the right from $t = 1$. Compare Fig. 10, where the ratio of the envelopes of the blue and red curves in the lower-right image is closely analogous to the blue curve here. The 1-norms of the approximation errors over $[-2, -1]$ are indicated. The slight irregularities at the left are the result of rounding error

The switch from the ∞- to the 1-norm changes the problem profoundly. The question becomes, how fast can φ be approximated by rational functions on $[-1, 1]$ in the 1-norm?

As we did with Fig. 10, let us get some insight by looking at the details of the approximation problem. The rational function (34) for the tanh rule is

$$r(t) = h \sum_{k=-(n-1)/2}^{(n-1)/2} \frac{\text{sech}^2(kh)}{t - \tanh(kh)}, \tag{38}$$

and for the tanh-sinh rule, it is

$$r(t) = h \sum_{k=-(n-1)/2}^{(n-1)/2} \frac{\frac{\pi}{2} \cosh(kh)\text{sech}^2(\frac{\pi}{2} \sinh(kh))}{t - \tanh(\frac{\pi}{2} \sinh(kh))}. \tag{39}$$

Figure 14 plots $|\varphi(t) - r(t)|$ for these two approximations.

For tanh quadrature, we know that $|\varphi(t) - r(t)|$ must diverge to ∞ as $t \to -1$ because of the log singularity of φ at $t = -1$. Yet the singularity is so weak that the divergence only shows up as a gentle upward drift in the blue curve at the left. Over the main part of the plot, $|\varphi(t) - r(t)|$ decreases steadily down to around 10^{-7}. The 1-norm, measured here over $t \in [-2, -1]$, is consequently very small, confirming via (37) the high accuracy of this quadrature rule. As $n \to \infty$, this 1-norm decays root-exponentially.

For tanh-sinh quadrature, again no approximation of φ is possible in the ∞-norm. In the 1-norm, however, one might expect that the convergence will now be almost-exponential. Indeed, $|\varphi - r|$ decays almost-exponentially as $n \to \infty$ over any domain bounded away from the singularity. But the 1-norm decay over the whole interval is in fact just root-exponential, as is suggested by the number listed being barely smaller than before. The following reasoning suggests why this must be. Since the essence of the matter concerns just a single log singularity, consider approximation of $f(x) = \log x$ on $[0, 1]$. Suppose rational approximations existed with faster than root-exponential convergence in the 1-norm. Then by integrating, we would get rational approximations to $g(x) = x \log x - x$ with faster than root-exponential convergence in the ∞-norm, which would contradict the evidence of Fig. 2.

If $\|\varphi - r\|_1$ decreases only root-exponentially as $n \to \infty$, how does the quadrature formula converge almost-exponentially? It appears that this depends on additional properties that go beyond rational approximation, involving analytic continuation of the integrand onto an infinitely-sheeted Riemann surface in exponentially small neighborhoods of the endpoints [50,54].

There remains the phenomenon of tapered exponential clustering, so vividly evident in Fig. 13. We do not yet have an explanation for this, nor a view of whether an approximate $\sqrt{k}$ dependence is genuine or just an artifact. This is a topic for ongoing research, where it would be good to investigate also the distributions of exponentially clustered nodes, also apparently tapered, that arise with the "universal quadrature" formulas of Bremer et al. [5,38].

7 Discussion

Exponential clustering of poles at singularities has been part of the landscape of rational approximation for half a century, though not many works in the area focus on this effect. (A fascinating previous example is [20].) Our motivation is that this clustering is what makes rational approximations so powerful, and understanding it enables one to improve existing numerical algorithms and develop new ones. We find these phenomena fascinating, especially the tapered clustering effect, and discovering that tapering also appears in double exponential quadrature was a bonus. The elucidation of these matters with the help of a sometimes seemingly endless program of numerical experiments will forever be associated in our minds with the Covid-19 shutdowns of 2020.

Here are some details of our computations. Figures 1, 2, 6 and 10 made use of the Chebfun `minimax` command [8], principally due to Silviu Filip, and Filip also kindly provided us with a modified code for the weighted minimax approximations of Figs. 2 and 6. For successful results in some of these problems, we applied a Möbius transformation of $[0, 1]$ to itself to weaken the singularity while preserving the space R_n. For the approximations of Figs. 2 and 6 on a complex disk, the AAA-Lawson algorithm was used as implemented in Chebfun [7,28], again with a Möbius transformation. Figure 3 was produced with the `confmap` code available at [59], which in turn calls `aaa` from Chebfun [27] and `laplace` from [59]. The `aaa` code was also used directly in Figs. 1 and 6, and `laplace` in Fig. 5. The Stokes and Helmholtz results of Fig. 5

were produced by experimental codes that are not yet publicly available developed with Abi Gopal and Pablo Brubeck, respectively. In Fig. 12, a least-squares problem was extended by a Lawson iteration (iteratively reweighted least-squares) to compute minimax approximations with preassigned poles. All the remaining results are based on straightforward computations in MATLAB and Chebfun.

References

1. Bailey, D.H., Borwein, J.: Hand-to-hand combat with thousand-digit integrals. J. Comput. Sci. **3**, 77–86 (2012)
2. Beckermann, B., Townsend, A.: Bounds on the singular values of matrices with displacement structure. SIAM Rev. **61**, 319–344 (2019)
3. Bernstein, S.N.: Sur l'ordre de la meilleure approximation des fonctions continues par des polynômes de degré donné. Mem. Acad. R. Belg. **13**, 1–104 (1912)
4. Bernstein, S.N.: Sur la meilleure approximation de |x| par des polynômes de degrés donnés. Acta Math. **37**, 1–57 (1914)
5. Bremer, J., Rokhlin, V., Sammis, I.: Universal quadratures for boundary integral equations on two-dimensional domains with corners. J. Comput. Phys. **229**, 8259–8280 (2010)
6. Brubeck, P.D., Trefethen, L.N.: Lightning Stokes Solver (manuscript in preparation) (2020)
7. Driscoll, T.A., Hale, N., Trefethen, L.N.: Chebfun Guide, Pafnuty Publications, Oxford. www.chebfun.org (2014)
8. Filip, S.-I., Nakatsukasa, Y., Trefethen, L.N., Beckermann, B.: Rational minimax approximation via adaptive barycentric representations. SIAM J. Sci. Comput. **40**, A2427–A2455 (2018)
9. Folland, G.B.: Introduction to Partial Differential Equations, 2nd edn. Princeton University Press, Princeton (1995)
10. Gauss, C.F.: Methodus nova integralium valores per approximationem inveniendi. Comment. Soc. Reg. Scient. Gotting. Recent. **39–76** (1814)
11. Gončar, A.A.: Estimates of the growth of rational functions and some of their applications. Math. USSR-Sbornik **72**, 445–456 (1967)
12. Gončar, A.A.: On the rapidity of rational approximation of continuous functions with characteristic singularities. Math. USSR-Sbornik **2**, 561–568 (1967)
13. Gončar, A.A.: Zolotarev problems connected with rational functions. Math. USSR-Sbornik **7**, 623–635 (1969)
14. Gopal, A., Trefethen, L.N.: Representation of conformal maps by rational functions. Numer. Math. **142**, 359–382 (2019)
15. Gopal, A., Trefethen, L.N.: Solving Laplace problems with corner singularities via rational functions. SIAM J. Numer. Anal. **57**, 2074–2094 (2019)
16. Gopal, A., Trefethen, L.N.: New Laplace and Helmholtz solvers. Proc. Nat. Acad. Sci. **116**, 10223 (2019)
17. Guo, B., Babuška.: The h-p version of the finite element method, part 2: general results and applications. Comput. Mech. **1**, 203–220 (1986)
18. Gutknecht, M.H., Trefethen, L.N.: Nonuniqueness of best rational Chebyshev approximations on the unit disk. J. Approx. Theor. **39**, 275–288 (1983)
19. Haber, S.: The tanh rule for numerical integration. SIAM J. Numer. Anal. **14**, 668–685 (1977)
20. Ingerman, D., Druskin, V., Knizhnerman, L.: Optimal finite difference grids and rational approximations of the square root I. Elliptic problems. Commun. Pure Appl. Math. **53**, 1039–1066 (2000)
21. Iri, M., Moriguti, S., Takasawa, Y.: On a certain quadrature formula (Japanese). Kokyuroku RIMS, Kyoto U. **91**, 82–119 (1970)
22. Lehman, R.S.: Development of the mapping function at an analytic corner. Pac. J. Math. **7**, 1437–1449 (1957)

23. Levin, E., Saff, E.B.: Potential theoretic tools in polynomial and rational approximation. In: Harmonic Analysis and Rational Approximation, Springer, pp. 71–94 (2006)
24. Mori, M.: Discovery of the double exponential transformation and its developments. Publ. Res. Inst. Math. Sci. Kyoto (RIMS) **41**, 897–935 (2005)
25. Mori, M., Sugihara, M.: The double-exponential transformation in numerical analysis. J. Comput. Appl. Math. **127**, 287–296 (2001)
26. Nakatsukasa, Y., Freund, R.W.: Computing fundamental matrix decompositions accurately via the matrix sign function in two iterations: the power of Zolotarev's functions. SIAM Rev. **58**, 461–493 (2016)
27. Nakatsukasa, Y., Sète, O., Trefethen, L.N.: The AAA algorithm for rational approximation. SIAM J. Sci. Comput. **40**, A1494–A1522 (2018)
28. Nakatsukasa, Y., Trefethen, L.N.: An algorithm for real and complex rational minimax approximation. SIAM J. Sci. Comput. (to appear) (2021)
29. Newman, D.J.: Rational approximation to $|x|$. Mich. Math. J. **11**, 11–14 (1964)
30. Okayama, T., Matsuo, T., Sugihara, M.: Error estimates with explicit constants for Sinc approximation, Sinc quadrature and Sinc indefinite integration. Numer. Math. **124**, 361–394 (2013)
31. Rakhmanov, E.A.: The Gonchar–Stahl ρ^2-theorem and associated directions in the theory of rational approximations of analytic functions. Math. Sbornik **207**, 57–90 (2016)
32. Ransford, T.: Potential Theory in the Complex Plane. CUP, Cambridge (1995)
33. Richardson, M., Trefethen, L.N.: A sinc function analogue of Chebfun. SIAM J. Sci. Comput. **33**, 2519–2535 (2011)
34. Rubin, D., Townsend, A., Wilber, H.: Bounding Zolotarev numbers using Faber rational functions. arXiv:1911.11882v2 (2020)
35. Runge, C.: Zur Theorie der eindeutigen Analytischen Functionen. Acta Math. **6**, 229–244 (1885)
36. Saff, E.B., Totik, V.: Logarithmic Potentials with External Fields. Springer, Berlin (1997)
37. Schwab, Ch.: p- and hp-Finite Element Methods. Clarendon Press, Oxford (1999)
38. Serkh, K.: Personal communication (September 2018)
39. Stahl, H.: General convergence results for rational approximants. In: Chui, C.K., et al. (eds.) Approximation Theory VI, pp. 605–634. Academic Press, New York (1989)
40. Stahl, H.: Best uniform rational approximation of $|x|$ on $[-1, 1]$. Mat. Sb. **183**, 85–118 (1992)
41. Stahl, H.: Best uniform rational approximation of x^α on [0,1]. Bull. Am. Math. Soc. **28**, 116–122 (1993)
42. Stahl, H.: Poles and zeros of best rational approximants of $|x|$. Constr. Approx. **10**, 469–522 (1994)
43. Stahl, H.R.: Best uniform rational approximation of x^α on [0, 1]. Acta Math. **190**, 241–306 (2003)
44. Stahl, H.R.: Sets of minimal capacity and extremal domains. arXiv:1205.3811 (2012)
45. Starke, G.: Fejér–Walsh points for rational approximation and their use in the ADI iterative method. J. Comput. Appl. Math. **46**, 129–141 (1993)
46. Stenger, F.: Numerical methods based on Whittaker cardinal, or sinc functions. SIAM Rev. **23**, 165–224 (1981)
47. Stenger, F.: Explicit nearly optimal linear rational pproximation with preassigned poles. Math. Comput. **47**, 225–252 (1986)
48. Stenger, F.: Numerical Methods based on Sinc and Analytic Functions. Springer, Berlin (1993)
49. Suetin, S.P.: Distribution of the zeros of Padé polynomials and analytic continuation. Russ. Math. Surv. **70**, 5–121 (2015)
50. Sugihara, M.: Optimality of the double exponential formula–functional analysis approach. Numer. Math. **75**, 379–395 (1997)
51. Takahasi, H., Mori, M.: Estimation of errors in the numerical quadrature of analytic functions. Appl. Anal. **1**, 201–229 (1971)
52. Takahasi, H., Mori, M.: Quadrature formulas obtained by variable transformation. Numer. Math. **21**, 206–219 (1973)
53. Takahasi, H., Mori, M.: Double exponential formulas for numerical integration. Publ. RIMS Kyoto **9**, 721–741 (1974)
54. Tanaka, K., Sugihara, M., Murota, K., Mori, M.: Function classes for double exponential integration formulas. Numer. Math. **111**, 631–655 (2009)
55. Trefethen, L.N.: Spectral Methods in Matlab. SIAM, Philadelphia (2000)
56. Trefethen, L.N.: Is Gauss quadrature better than Clenshaw–Curtis? SIAM Rev. **50**, 67–87 (2008)
57. Trefethen, L.N.: Series solution of Laplace problems. ANZIAM J. **60**, 1–26 (2018)

58. Trefethen, L.N.: Approximation Theory and Approximation Practice, Extended edn. SIAM, Philipedia (2019)
59. Trefethen, L.N.: Lightning Laplace Software. https://people.maths.ox.ac.uk/trefethen/lightning (2020)
60. Trefethen, L.N.: Numerical conformal mapping with rational functions. Comput. Methods Funct. Thy. **20**, 369–387 (2020)
61. Trefethen, L.N., Weideman, J.A.C.: The exponentially convergent trapezoidal rule. SIAM Rev. **56**, 385–458 (2014)
62. Vyacheslavov, N.S.: On the uniform approximation of $|x|$ by rational functions. Sov. Math. Dokl. **16**, 100–104 (1975)
63. Walsh, J.L.: Interpolation and approximation by rational functions in the complex domain. Am. Math. Soc. (1935)
64. Wasow, W.: Asymptotic development of the solution of Dirichlet's problem at analytic corners. Duke Math. J. **24**, 47–56 (1957)
65. Wegert, E.: Visual Complex Functions: An Introduction with Phase Portraits. Birkhäuser, Basel (2012)
66. Widder, D.V.: Functions harmonic in a strip. Proc. Am. Math. Soc. **12**, 67–71 (1961)
67. Xie, T.F., Zhou, S.P.: The asymptotic property of approximation to $|x|$ by Newman's rational operators. Acta Math. Hungar. **103**, 313–319 (2004)
68. Zolotarev, E.I.: Application of elliptic functions to questions of functions deviating least and most from zero. Zap. Imp. Akad. Nauk. St. Petersburg **30**, 1–59 (1877)